GODDARD SPACE FLIGHT CENTER

Advanced Weather Prediction Technologies: NASA’s Contribution to the Operational Agencies

Gap Analysis Appendix

May 31, 2002

Greenbelt, Md. 20771

1. Introduction

This appendix describes the gap analysis performed in support of the notional vision architecture described in the study report. The architecture was used as a starting point to perform these gap analyses.

The first portion of the appendix is the matrix of possible technologies that the study team determined to be candidates for the 2025 architecture. The matrix identifies the anticipated maturity levels, and the analyses were performed for key areas where significant gaps were flagged.

The second portion of the appendix contains the details of the gap analyses.

Notes:

This appendix is not meant to stand alone from the study report; very little explanation of the technologies' roles in the architecture is given. To get the complete picture this appendix must be read after the study report, with frequent references to the study report likely.

2. Summary of Potential Technology Gap Areas

The table in Attachment 1 lists the technologies identified as possible candidates to be used in the notional vision architecture described in the main body of the Weather Prediction Technology Investment Study report. It is structured to follow the top level of the architecture, namely the Observing Systems, expansions on that Observing System, the Model and Data Assimilation System (M/DAS), and the Knowledge Delivery system. The next level down in the chart is organized by function or subsystem, with the next levels being technologies within that function or subsystem.

This chart represents the best estimates of the team, but was not exhaustively worked for completeness of technologies listed nor their anticipated readiness levels. It was not meant to be as formal as TRLs but rather as a starting point for future, in-depth studies. The levels of anticipated readiness are described at the end and are listed by categories across the top. The engineers and technologists on the study team researched technologies in their area of expertise and extrapolated developing technologies to the 2025 timeframe, a difficult task given its highly subjective nature.

The column definitions are as follows:

- Performance Objective – what the technology needs to be able to provide in 2025; wherever possible a number was developed, but that was not always the case given the nature of advanced concepts.
- Capability Maturity – current status
- Challenges and Investments – given investment levels, how likely is it that this technology will be ready to meet the performance objectives in 2025? The higher the number, the less likely, with red being the color that identifies a gap.
- Technology columns – this identifies the technology nature of the gaps
- Science columns – this identifies whether a science gap is in the area of research or infrastructure

3. Observing System Gap Descriptions

As stated in the main body of the report, the success of a weather forecast depends heavily on how well the initial conditions are portrayed. The "goodness" of this portrayal is determined both by the accuracy and by the "representativeness" of the measurements as well as their timeliness. This section discusses the technology gaps in both the means to collect representative measurements as well as the means deliver them to their destination in a timely fashion.

Collection technologies include:

- Sensor technology needed to measure the various parameters
- On-board computing needs to process the data
- Guidance and navigation technologies needed to precisely determine the collection location
- Collection management and control to orchestrate the data collection

Delivery technologies primarily involve the communications and networking capabilities needed to get the collected data to the users in the required timelines.

3.1. Wind Remote Sensing - Anticipated Wind Lidar Technology Needs

The assessment of technology requirements for a wind lidar is complicated because there are two techniques proposed for making this measurement. The basic idea of both is to measure the Doppler shift of light scattered by molecules and/or particles carried by the wind. The direct detection Doppler lidar method uses a high spectral resolution optical filter (often a Fabry-Perot interferometer) to measure this shift using the atmospheric backscattered laser energy from either molecules or aerosols. The coherent or heterodyne Doppler lidar converts laser light backscattered from aerosols or clouds from optical to radio frequencies and uses RF spectral analysis techniques to measure this frequency shift. Current direct detection approaches use near UV wavelengths (λ~350 nm) for the molecular Doppler wind measurements while the heterodyne technique proposes to operate in the near IR (λ-2 microns). Both of these approaches have been demonstrated using ground based lidars and have been studied extensively for spaceborne operations.

In general the capability to measure winds of both lidar approaches is a function of the signal-to-noise ratio (SNR) of the signal detected from the atmosphere. In both cases the SNR will be a function of instrumental characteristics (e.g. laser energy and repetition rate; number of laser shots averaged; telescope collection area; detector quantum efficiency; optical throughput), spacecraft related characteristics (e.g. orbital height; nadir angle; pointing accuracy and control) and atmospheric effects (e.g. spatial distribution (horizontal and vertical) of the target particles (aerosols and/or molecules); wavelength dependent molecular and aerosol backscatter coefficient; two way atmospheric transmission; cloud distribution, height and optical properties). It becomes clear that any detailed analysis of the technology trade space will be highly dependent on the specifics of the implementation. To complicate this further, the details of how the SNR relates to the characteristics of desired wind product (accuracy, vertical and horizontal resolution) are different for coherent detection and direct detection Doppler lidars. Finally, the scaling of technologies for the individual approaches to larger sizes (with thereby improved capabilities) is not directly comparable. For example the heterodyne approach requires a telescope with diffraction limited performance while the direct detection approach can use a much lower quality telescope. On the other hand, heterodyne detection has high out of band noise rejection and so will be immune to solar background noise even in daylight while direct detection signals must be determined in the presence of background noise. In any case the measurement by either technique is extremely challenging.

Fortunately, recent engineering studies of both approaches have established reference baselines for both coherent and direct detection approaches, and for the purposes of evaluating technology needs and gap analysis we can use the results of those studies to extrapolate the needs of the future. Consider the following analysis:

The measurement requirements proposed for this study are as follows:

1. Global measurement of 2-D winds with precision of 1 m/sec.
2. Horizontal resolution 25 km x 25 km.
3. Vertical resolution 0.25 km
4. Temporal resolution 3 hours.

We define a "measurement" to be an altitude profile measured by the lidar viewing in a single direction. This means to get a 2-D wind determination in a single horizontal resolution element, but at all required altitudes, requires 2 "measurements." We calculate the number of "measurements" required per day as follows:

Area of a resolution element is 25 x 25 = 625 km^2.

Area of the Earth is $4\pi r^2$ where $r = 6.36 \times 10^3$ km. This equals 5.1×10^8 km^2.

Therefore, the number of horizontal resolutions elements on the earth is about 816,000.

Each of these elements must be visited twice every 3 hours to yield the required 2-D temporal resolution, so the total number of daily altitude profiles is

$$816{,}000 \times 16 = 1.3 \times 10^7 \text{ per day.}$$

A recent engineering study of wind lidar capabilities presumed that a single lidar had the ability to make about eight line-of-sight wind measurements per minute. If that lidar system is baselined then it could make

8 x 24 x 60 = 11520 line of sight measurements per day or

4 x 24 x 60 = 5760 horizontal wind measurements per day

It would therefore require roughly 2200 such lidars operating continuously to meet the horizontal wind measurement requirements.

This factor of 2200 could be made up in various ways (eg. 22 platforms with lidars of 100 times more capability or say 100 platforms having lidars with 22 times greater capacity.)

However, in the engineering studies the lidar system only had the ability to resolve vertically at 1 km resolution and had a precision of 3 m/sec rather than 1 m/sec. Scaling up the performance of the lidar to meet the more stringent requirements assumed in this study implies that we need about another order of magnitude improvement in the lidar capability above the 2200 already discussed.

If we assume that 100 platforms could be utilized, then we need to achieve about a factor of 100 improvement in sensitivity of the lidar over the next 25 years. Improving lidar sensitivity can come about in several ways:

1. Increase in laser power (pulse energy X repetition rate).
2. Increase in the collection area of the receiver optical system.
3. Increase in the efficiency (quantum efficiency or throughput) of the detection system.

Detector quantum efficiency is already relatively high for both the near IR coherent and near UV direct detection systems. It might be possible to achieve some level of improvement, although it is difficult to imagine that this increase could be more than about a factor of 2.

Current telescope systems for spaceborne lidars have an aperture of about 1 meter. If this could be scaled up by a factor of 3 then we would have found about an order of magnitude in sensitivity, as the SNR scales as the area. (Note: Background noise would also increase for a direct detection lidar and pointing knowledge and control requirements would increase for a coherent system). A premium here is placed on increasing aperture without significantly increasing mass. In addition, single satellite lidar systems may be required to obtain multiple perspectives by slewing the FOV to different azimuth angles by rotating the telescope or an external scanning optic.

The remaining factor of 5-10 would require improvements in many areas. Among these are the following:

- Laser efficiency (conversion of electrical power from spacecraft into light energy).
- Solar power conversion and storage efficiency (for high power drain lasers on spacecraft).
- Pointing system performance and efficiency.
- Frequency conversion efficiency.

3.2. Microwave Remote Sensing Technology

Microwave measurements have become an important input into the weather forecast models of today. Many of these measurements will need to be performed on a daily or even hourly basis on a global scale. In particular this section will try to address some of the technologies, which will allow measurement of these parameters on a global scale from space.

There are several passive and active microwave measurements which are useful for weather forecasting. This section will attempt to detail the technology improvements necessary to meet the requirements of this future forecasting study. Measurements that are particularly well suited to microwave remote sensing include soil moisture (and other similar surface parameters), atmospheric temperature, and atmospheric moisture.

3.2.1. Soil Moisture

Currently the 6.9 GHz radiometer on AMSR is the only high-resolution passive microwave measurement of soil moisture available. This frequency is not the optimum for measuring soil moisture. A better choice would be the 1.4 GHz remote sensing band. This frequency is much more sensitive to soil moisture and can penetrate through more vegetation than higher frequencies. The difficulty with using 1.4 GHz to measure soil moisture is that the wavelength is 5 times longer than 6.9. This results in an antenna whose diameter is 5 times as large. The most recent proposals for measuring soil moisture from space would produce resolutions on the order of 40 km. These missions have antennas that are about 6m in diameter. To move to a measurement of 1 km would require increasing the antenna size to approximately 240 m. This seams like a tall order since that antenna would also have to spin at about 120 RPM (2 Hz). Given these assumptions, potential solutions can be imagined.

3.2.1.1. High Resolution Array Technology

The technique of Synthetic Thinned Array Radiometry (STAR) similar to what was proposed for HYDROSTAR or the European Space Agencies' Soil Moisture Ocean Salinity (SMOS) missions can be scaled to higher resolutions than are being proposed today. Within 10 years at current levels of funding a 10 km soil moisture measurement could be implemented. The configurations that have been proposed are probably limited to 10 km resolution. Above this resolution the system becomes so large that decorrelation becomes a problem for the available bandwidth. The alternative is something that has been called the Doppler radiometer. It is an interferometer made of three or more radiometers, which are phase locked to each other. These radiometers would fly in a formation separated by the maximum diameter of the antenna required to achieve the desired resolution. A great deal of further study is needed to determine if this configuration can actually achieve the sort of sensitivity that is useful for a soil moisture measurement. A modest increase in funds for the study of this concept and methods of formation flying could lead to a space demonstration within 10 to 15 years.

3.2.1.2. Active/Passive Combined Algorithm

An alternative approach is one that has been suggested by Ulaby and others. This involves a combination of an active and passive measurement of soil moisture. The advantage of this approach is that it potentially takes advantage of the strengths of each measurement. The passive measurement would provide the high accuracy low-resolution soil moisture to use as a reference for the less accurate but potentially very high-resolution active measurement. One could envision a measurement, which included a 10 km STAR radiometer imaging primarily on each side of the spacecraft and an unfocused Synthetic Aperture Radar (SAR) with 1 km spatial resolution. This set of measurements would then be combined in a statistical way to provide a measurement, which has 1 km resolution and high sensitivity to soil moisture in the presence of vegetation. It is again unclear if this method will ever be useful. It currently has not been demonstrated in any field experiment. If it can be accomplished, the development of a field instrument and funding for the development of the algorithm will be required to prove that it can work.

3.2.1.3. Large Real Aperture Approach

Perhaps the least elegant but simplest approach to the electrical design is the large single aperture. This aperture would have to grow to several hundred meters to make a 1 km measurement possible. This certainly seems to be unlikely to happen in the immediate short term. The technology that would make something like this possible is thin film inflatable antennas. This technology can produce very lightweight antenna structures with reasonable antenna characteristics.

3.2.2. Atmospheric Temperature and Moisture

Currently, atmospheric temperature and moisture are measured reasonably well, but at course horizontal and vertical resolutions. The frequencies of interest for these parameters typically range from approximately 19 – 85 GHz. These higher frequencies will require much smaller antenna systems than those for measurement of soil moisture. In all likelihood, there will not be a technology gap for these measurements.

3.3. On-Board Processing Technology

As our concept of a SensorWeb evolves, more of the processing needed to support its intelligence will need to be moved from the ground to space platforms. Some of this processing, such as data calibration and reduction, will be relatively simple (in terms of computing costs) while other functions, such as automated event recognition to enable SensorWeb reconfiguration, may be somewhat computationally expensive.

3.3.1. Anticipated On-Board Processing Needs

Although it is impossible to estimate the exact needs without a full concept development for how the SensorWeb will operate, we can at least make some broad generalizations. Data correction and reduction are currently done on workstation-class computers with capabilities in the low 100's of MFLOPS range. Other functions, such as initial forays into event detection, are done on higher-end workstation- or mainframe-class computers with capabilities in the high 100's of MFLOPS to GFLOPS range in a research mode. The anticipated needs in an operational mode are not yet known.

3.3.2. Anticipated On-Board Processing Capabilities

Currently, NASA's most powerful radiation-hardened computer matches the capability of a 80486 processor. However, non-hardened, higher end computer processors could be flown in space using various software and hardware techniques to overcome radiation effects. Other technologies being flown commercially are already putting high-end computing capabilities into space.

Boeing's recently launched Thuraya communications satellite has a digital signal processing power equivalent to 3000 Pentium III computers. Unfortunately, the power requirements of this system necessitate solar panels nearly 35 meters across for electrical power generation and a 7.4 square meter radiator to dissipate the heat, certainly beyond practical limits for a constellation of LEO platforms.

If a practical application of these computing capabilities is to come about, especially for a constellation of LEO platforms, a means of reducing the power requirements is needed. Studies by the Space Telescope Science Institute and the Jet Propulsion Laboratory (JPL) are examining the potential for high performance computing in space with lower power requirements. Experiments underway are running prototype systems capable of 30 MFLOPS per watt. It is hoped that computing capabilities of 300 – 1000 MFLOPS per watt will be available to support the Next Generation Space Telescope expected to fly as early as 2008.

If Moore's Law is applied to these numbers, we can estimate the expected on-board computing capability over time. Figure 3-2 shows this projection of on-board computing capabilities through 2025. The lower line (in red) shows the expected growth of computing capability per watt of power. Because systems of 20 such processors are envisioned, the computing capability of a 20-node system is shown by the upper line (in green). Thus, it seems reasonable that an on-board computing capability of 1000 GFLOPS – or just a little bit better than today's ground-based super computers – can be expected by 2025.

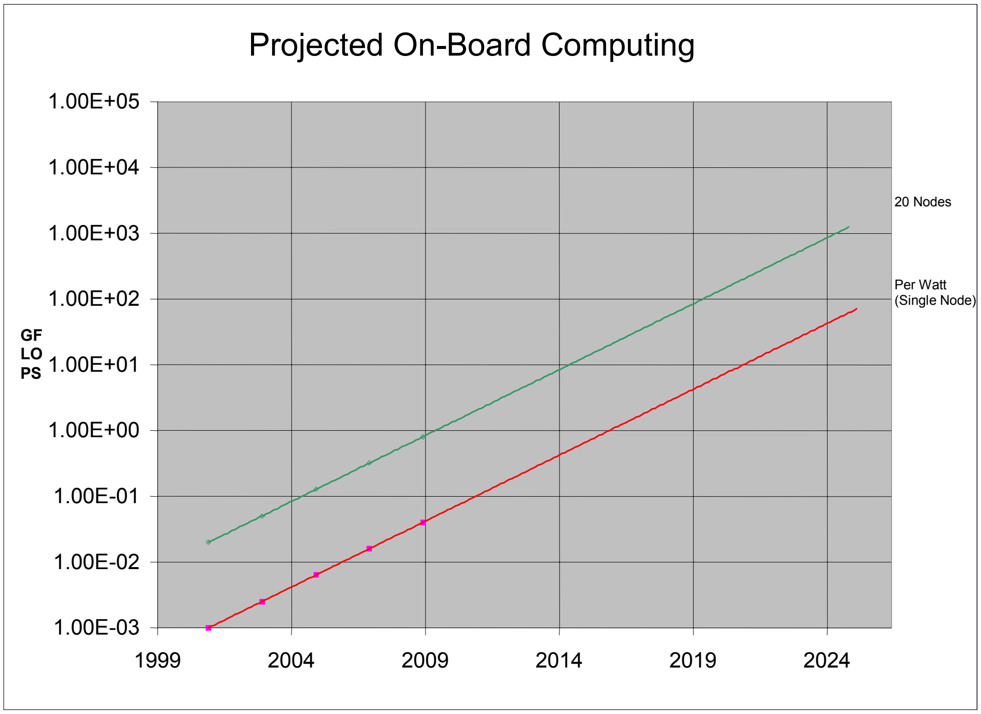

Figure 3-1 Projected On-Board Computing Capabilities

3.3.3. Gap Analysis

It is difficult to assess whether a technology gap exists based on the lack of a firm concept of operations. However, based on the initial concepts outlined in the study report (where only data calibration/reduction and initial QC are accomplished on board), it appears as if there will be no gap in on-board computing capabilities.

However, event detection and recognition algorithms (whose computational complexity is not yet known) might tax the expected capabilities. Further studies into potential science applications and their computational costs are needed to fully understand these needs.

3.3.3.1. Technology Shortfalls

As Boeing's Thuraya communications satellite demonstrates, high performance computing in space is possible, even with today's technology. However, it comes at great costs in terms of weight, power, and thermal considerations. If these computing capabilities are to become a possibility for a LEO constellation, the computing capability per watt is a critical factor. Although current prototype efforts are promising, it remains to be seen if these technologies will be scalable for future needs.

3.3.3.2. Trade Areas

The key trade area to be investigated relates to functions to be processed on board vs. on the ground. These trades must be weighed against how much communications bandwidth is expected to be available. If greater communications capacity can bring data to the ground in near real time, then processing can be done on the ground, reducing the need for on-board processing requirements. However, if communications bandwidth is limited, or if the science of event detection and recognition dictates an immediate response, then high performance computing in space must be considered.

If the vision of low-power computing does not come to fruition, then trades must be made between computing capability and power/thermal considerations.

3.3.3.3. Future Technology Needs

The future weather architecture outlined in the study report will certainly require increased computational resources on-board the space platforms. With the power requirements of the instruments, especially the active sensors, power considerations will be a limiting factor. Thus, increased computing capability with reduced power costs is going to be crucial. If these power reductions cannot be realized, increases in power generation capabilities (e.g., more efficient solar panels) and better thermal management will be a must.

3.3.3.4. Recommendations

As in the ground-based computing portion of this study, we must keep an eye on the computer industry. Furthermore, we must maintain an open dialog with the computer research community (such as the REE project) so they remain aware of our future computational needs.

ESTO should also support research into developing more efficient computational systems and algorithms to make better use of the available computational resources and support research into smarter analysis and forecast algorithms.

ESTO should support a follow-on effort to flesh out a concept of operations in order to more fully identify what processing requirements are needed in space. In the current version of the notional architecture, only minimal data processing is accomplished on board the spacecraft conducting remote sensing measurements. It is conceivable that some portion of the data processing should be done on the spacecraft that would benefit either the efficiency of the system or quality of the collected data and forecast products.

3.4. Guidance, Navigation, and Control (GN&C)

In addition to the sensor and computing technologies needed, improvements in spacecraft GN&C will also be needed. As higher resolution measurements are made, it is increasingly more important to have a better understanding of exactly where the measurements are being made to ensure that they are representative of the true state of the atmosphere. With the larger number of space platforms envisioned, the cost and complexity of managing the operation and control of the constellation could become prohibitive. The next sections discuss the technologies needed to address these concerns.

3.4.1. Global Positioning System (GPS) Navigation

Today's GPS can provide all the LEO weather satellites with latitude, longitude and altitude to an accuracy that is better than the minimum needed to locate the data collected.

This is not true for the GEO satellites. These satellites are above the GPS constellation and can only receive GPS signals from those GPS satellites that are close to setting (or rising) behind the earth. There is every theoretical reason to believe that satellites in GEO orbit will be able to use these signals to achieve adequate positional accuracy, but so far, no mission has demonstrated this. The GEO weather satellites will need this capability because accurate satellite location is needed for accurate, automatic location of images and other data produced by these satellites.

No current efforts are planned to explore this technology. Figure 3-2 shows the relative GEO GPS levels of performance (or technology readiness) and the approximate time each level could be obtained. Trades on performance (and benefits) of obtaining this capability compared to their costs must be conducted to determine if the appropriate performance level can be reached.

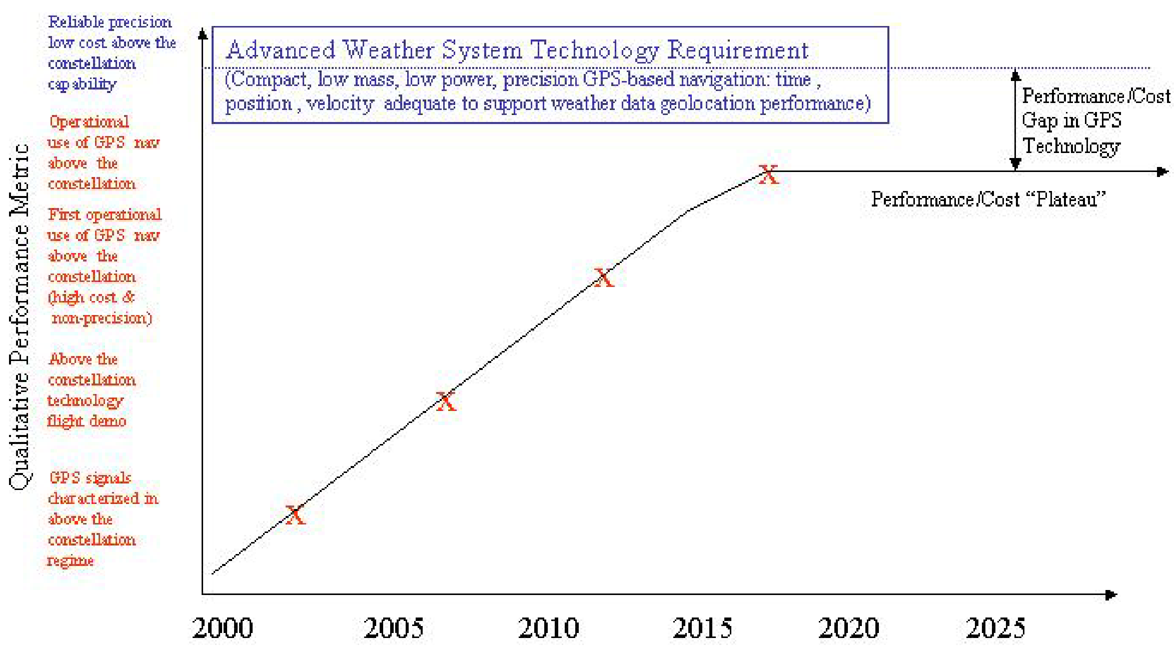

Figure 3-2 GPS Navigation Gap Analysis

3.4.2. Drag-free Control

In typical ESE missions significant ground resources are applied to the orbit determination and orbit prediction tasks. It is well known that the most uncertain part of solving the Low Earth orbiting flight dynamics problem is the prediction of atmospheric drag levels. The drag force from the Earth's atmosphere not only tends to decay spacecraft orbits, but also can vary significantly from day to day. This uncertainty drives the need for increased spacecraft tracking and detailed orbit modeling, determination and analysis to measure the vehicle positions as well as control them with periodic propulsive maneuvers.

The effective elimination of drag from the spacecraft flight dynamics problem reduces any orbital decay to purely gravitational perturbations which are known quite well and which can be compensated for with appropriate analysis. Through the use of an integrated accelerometer package on a spacecraft, most likely consisting of a floating proof mass in an internal chamber with electrostatic (capacitive) sensing and actuation, a high specific impulse (Isp) thruster, and a low-cost processor with appropriate filtering/control algorithms, a closed loop drag free control system can be synthesized. Such a drag free system will:

- Eliminate the effect of drag on each spacecraft to prevent decay of the orbit (using virtually insignificant continuous and non-interfering thrust).
- Continuously maintain the constellation elements within their boxes to avoid undesirable interactions
- Avert the need to shut down the mission every 1-4 weeks to perform a delta-vee orbit correction.
- Maintain precise knowledge of the orbital position of the vehicles continuously without sensitivity to upsets, bit-flips, etc. and without the need for expensive sensors. This will enable vastly improved geolocation performance to enable us to meet specs for such tasks as wind speed measurement.
- Avoid the need for complex algorithms for collision avoidance and large scale constellation maintenance.

Figure 3-3 shows the expected performance capabilities and the requirements of the system outlined in the report. Two performance metrics – acceleration cancellation levels and the number of spacecraft for system level application – are presented over time. Acceleration cancellation technologies already planned are anticipated to meet the requirements of the system envisioned. The second metric, however, is dependent upon two technologies, (1) low-cost, moderate performance, drag free (floating proof-mass) sensor development and (2) algorithms which enable us to use this technology at the constellation level, rather than at the single spacecraft level. Currently, no systems are planned to implement these technologies in constellations with the number of spacecraft envisioned. Significant development in this area is required.

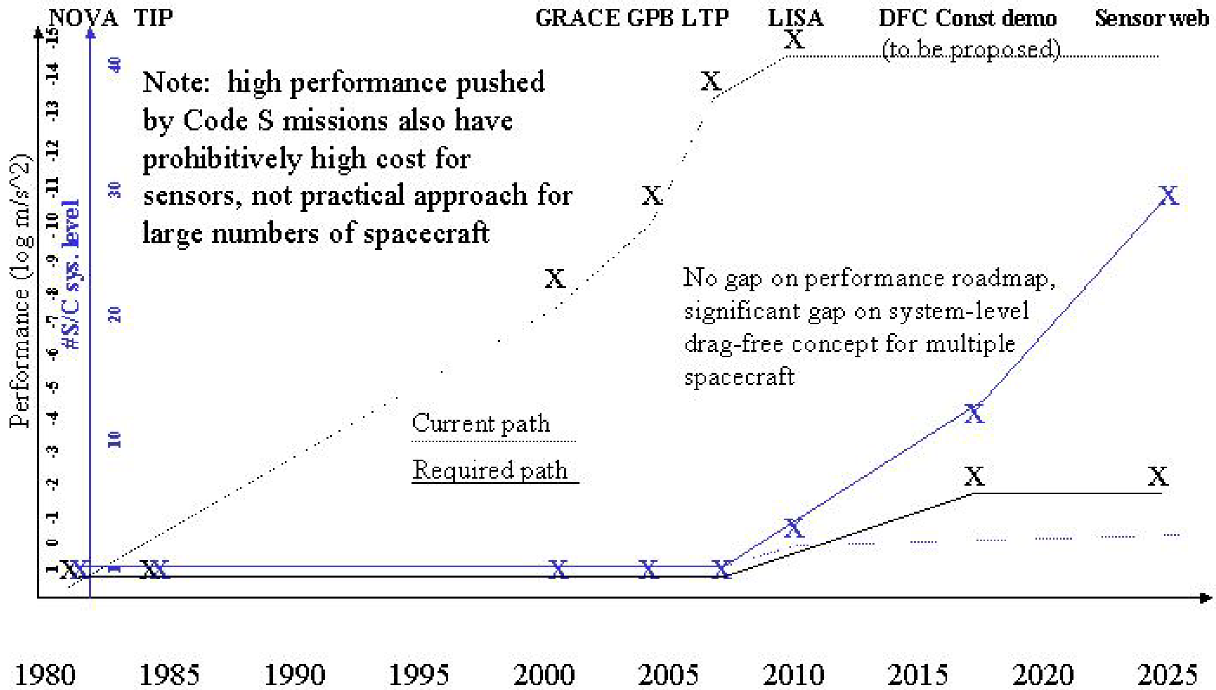

Figure 3-3 Drag Free Control Gap Analysis

3.5. SensorWeb Management/Control (SWM/C) Technology Gap

As envisioned, the SensorWeb will require rapid, nearly seamless communication between assets located in space, in the air, on the ground, and at sea. An overarching "intelligence", referred to here as the SensorWeb Management/Control (SWM/C), would manage the assets to make regularly scheduled data collections and to optimize the scientific targets of opportunity. The SWM/C would provide coordination between command/control for widely disparate collection platforms and complex dynamic planning and scheduling. Because of the unprecedented configuration of the proposed SensorWeb and the complexity of scheduling the assets technology gaps are evident.

3.5.1. SensorWeb Command/Control

The SensorWeb Command/Control will require all of the standard operating components seen in today's satellite systems. It will also require the services provided by such components to be extended across all observing systems assets (e.g. aircraft, ships, ground-based sensors, etc.). Such components include:

- Data architecture to identify the major components of the overall observing system
- System management architectures that provide for the organizational and management of the operations environment of the assets
- Control interfaces that provide a mechanism to operate and manage the assets
- Decision support components to operate the assets and process commands

3.5.1.1. SensorWeb Dynamic Planning and Scheduling

Planning and scheduling of the SensorWeb assets addresses the problem of formulating a sequence of commands that will result in achieving a desired scientific goal. A possible scenario for SensorWeb operations would be three “operating modes”:

- *Normal operating mode* would schedule regular collections of operations that would satisfy the basic requirements of the data assimilation system. This would include making measurements of temperature, moisture, wind, etc. at the appropriate temporal, spatial, and spectral resolutions. This mode would also address the scheduling of data points that must be re-sampled due to initial flagging by the meteorological quality control. The main driver of this mode of operations would be the data assimilation system.
- *Opportunistic science mode* would be used to capture events of specified scientific interest. The observing system elements and/or the data assimilation system would alert the planning/scheduling algorithm to perform intensive data collections focused on specific locations. For example, if a satellite detects conditions favorable for severe weather (perhaps by using on-board event-detection algorithms) the planning/scheduling component would interact with the data assimilation system to predict the location of the event over the next several hours and schedule high-resolution data collections accordingly. The entire lifecycle of the severe weather outbreak could then be captured.
- *Field experiment mode* would be used to manually select regions for intensive observations. This mode would be particularly useful for research studies that require higher resolution data over a specified location over a period of time.

It is unlikely the SensorWeb would ever operate in a single mode. Rather, to maximize the scientific benefit, an optimal combination of the three modes is necessary. Defining the optimal combination of such a dynamic system is a grand challenge of building the SensorWeb.

3.5.2. Anticipated Technology Capabilities

NASA and other government agencies are now formulating the roadmap to develop intelligent Distributed Spacecraft Systems (DSS). In recent studies, DSS is defined as a spatially distributed intelligent network of multiple space assets, collaborating as a collective unit, and exhibiting a common system-wide capability to accomplish shared objectives. This work is significant to the current study because there is considerable overlap in the command/control system requirements for DSS missions and for SensorWebs. Proposed Earth Science DSS missions that may require enhanced command/control capabilities over the next 10-15 years include:

- Global Precipitation Mission (currently scheduled for launch in 2007)
- Leonardo
- Topography and Surface Deformation
- GPS Atmospheric Sounding

There are now investigative activities underway to prepare for DSS missions. For example, two recent investments by the NASA Office of Aerospace Technology are funding studies into the development of discrete event controllers for autonomous, distributed spacecraft command/control, and for autonomous command/control for formation flying. There is also a significant amount of work underway in Space Sciences, with much of the research for distributed spacecraft problem solving being performed by the Jet Propulsion Laboratory. At Goddard Space Flight Center work is underway to develop the so-called "goal-oriented commanding" that is designed to perform high-level tasking of a constellation of satellites with a minimal amount of human intervention.

Although the DSS studies will likely provide some benefit in designing a SensorWeb command/control system, they are limited because they address only on-orbit assets. There are no studies currently investigating the design of a command/control system that manages a diverse suite of assets that would be needed by the proposed SensorWeb.

Figure 3-4 displays a likely technology capability roadmap for command/control capabilities of distributed spacecraft systems. The data contained in the figure is based upon an Earth Science Enterprise planning workshop conducted in 2001. The analysis suggests significant strides will be made in autonomous spacecraft control and scheduling over the next five years, with demonstrated mission capability likely by the end of the decade.

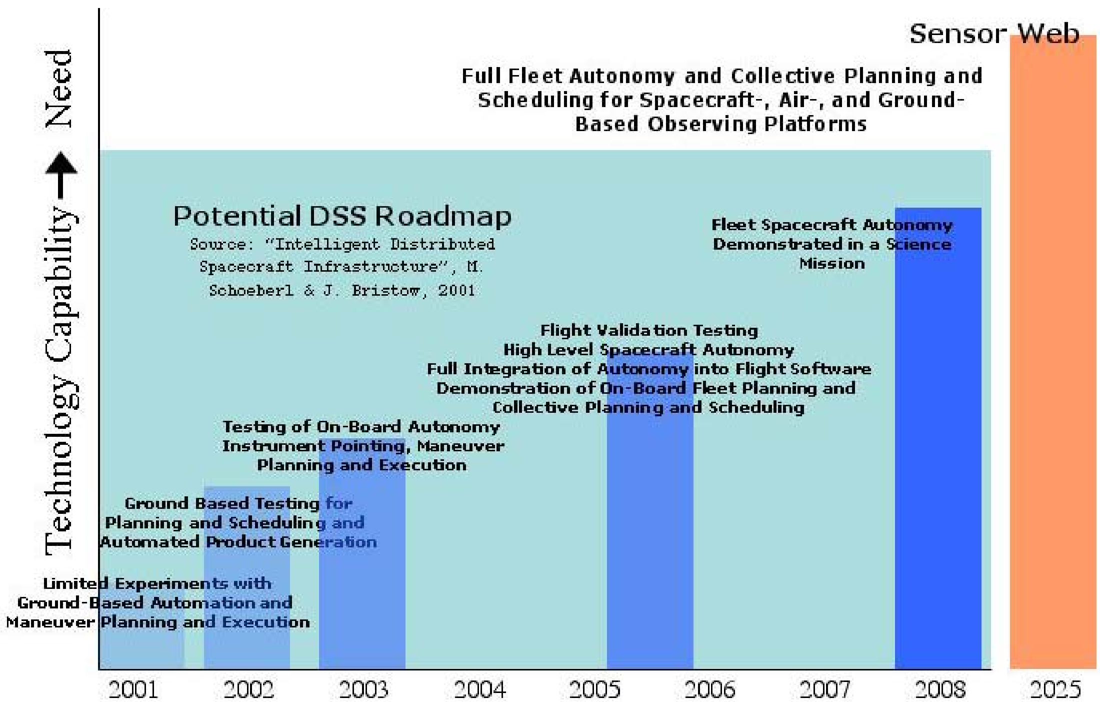

Figure 3-4 Anticipated technology roadmap for addressing command/control and planning/scheduling of DSS missions.

3.5.3. Gap Analysis

Although not entirely quantifiable, a technology gap appears to exist for the successful development of the SWM/C components of the SensorWeb. The gap is related to the diversity of the assets that must be managed, and the time constraints placed on the highly complex scheduling necessary.

3.5.3.1. Technology Shortfalls / Future Needs

Based upon the data contained in Figure 3-4, command/control and planning/scheduling algorithms for satellite constellations will likely be available by the end of the decade. Studies associated with the requirements analysis and the design of such algorithms will benefit the design of similar SensorWeb components. However there does not appear to be current or planned studies involved with linking services for satellite-based observing systems with air- and ground-based observation networks.

For the SWM/C, "asset awareness" and the complexity of the optimization problem appear to be the most significant gaps. In order to optimize data collection to maximize scientific return, the scheduling algorithm would be required to identify accurately the current and future locations of all assets, as well as deployment times and overall availability of rawinsondes, unattended aircraft, drifting buoys, ships, and current and anticipated states of many other resources. Therefore, the scheduling algorithm must have detailed, up-to-the-minute knowledge of perhaps tens of thousands of assets, and must perform scheduling decisions within a matter of seconds. Such decisions must be based upon weighing requests made under the three operating modes of the SensorWeb and rapidly formulating the "best" decision. Although similar algorithms exist today (goal-oriented commanding of spacecraft and even computer chess games that anticipate and score future moves of the chess pieces are relevant examples), the requirements of the SensorWeb require a significant step forward in both hardware and software technology.

3.5.3.2. Recommendations

Future studies should serve to bridge the apparent gap between relevant research on DSS and the requirements of the proposed SensorWeb, and should attempt to quantify at a low level the requirements for planning and scheduling within the SWM/C. Simulation of the SensorWeb environment would be one approach to understand the magnitude of the technology gap and to assist researchers in addressing the challenges presented by a complex observing system.

3.6. Communications Technology -- GEO Satellites

The GEO satellites would downlink their data directly to ground stations and the commands would be uplinked from the same stations. On the ground, commercial communications links would carry the data and the commands between the ground stations and the weather system command and control center(s). There are no technology concerns with any of these communications links.

3.7. Communications Technology -- LEO Satellites

3.7.1. Requirements that drive the communications System

A key requirement highlighted in the main report is a data latency of no greater than 15 minutes. This requirement drives the system architecture to use a space-based communications system similar to the present Tracking and Data Relay Satellite System (TDRSS). The capability of the satellites that compose the present TDRSS could do this mission.

Since the LEO weather satellites could have coverage at locations around the earth at a given time, a minimum space-based communication system requires three full-capability communications satellites located approximately 120 degrees apart. The ground terminals for each of these satellites must have adequate communications with the users of the weather data and the Command and Control Center. (Today's TDRSS has satellites located to the East and West of the United States with adequate capacity, but there is only one satellite with partial capability over the Indian Ocean. It is expected that by 2025, global TDRS capability will be available.)

3.7.2. Description of the Tracking and Data Relay Satellites (TDRSs)

The latest generation of TDRSs has three modes of receiving data from LEO satellites. They each include a Multiple Access (MA) service, a K-band Single Access (KSA) service and an S-band Single Access (SSA) service. The MA is an S-band phased array that is able to receive up to 3 Mbps from five LEO satellites simultaneously. Each TDRS has two Single Access Antennas and each antenna includes a KSA service that can receive up to 300 Mbps and a SSA service that can receive up to 6 Mbps. The KSA service can receive data from one LEO satellite and simultaneously can receive SSA data from the same satellite or from a second, nearby LEO satellite. The dead time between the end of one SA contact and the start of the next is typically 1.5 minutes for the newest TDRSs, but was only about 30 seconds for the original TDRSs. For this paper, it is assumed that future TDRSs will have dead times of less than 30 seconds.

3.7.3. Low Data Rate Spacecraft/Sensors (Multi-Access)

Table 3-1 gives the estimated data rates for the satellites that would make up the LEO constellation of weather satellites. It is seen that only the 8 LIDAR satellites have data rates suitable for the MA service. The three TDRSs would each handle 2 or 3 LIDAR satellites simultaneously. The data rate would be about 300 kbps or about 25% larger than the 241 kbps shown in the table. This is to allow for the dead time while the LEO satellite switch between one TDRS and another and to allow for the occasional data rate surges that are higher than average.

Table 3-1 Estimated LEO Weather Satellite Data Rates

		These instruments may be grouped on one satellite				LIDAR Satellite	RADAR Satellite
		EO Imager	MW Sounder	MW Imager	IR Sounder		
	Number of Instruments or Number of Satellites	13	13	13	13	8	13
	Time to Achieve Global Coverage hours	1	1	1	1	3	1
	Swath Width (km)	2700	2700	2700	2700	2700	2700
	Spatial Resolution (km) (nominal)	1	25	25	25	25	25
	Spatial Resolution (km) (surge)	0.5	5	5	10	10	1
	Period (min)	98	98	98	98	98	98
	Data size per observation (bits)	16	10	10	16	10	16
	Metadata / observation (bits)	1	40	40	512	512	40
NOMINAL	# Channels or Vertical Samples	12	7	10	200	320	20
	Observations / swath	2700	108	108	108	108	108
	Swaths / sec	6.808	0.272	0.272	0.272	0.272	0.272
	Nominal Data Rate Mbps	3.383	0.003	0.004	0.104	0.104	0.010
SURGE	# Channels or Vertical Samples	12	7	10	200	320	80
	Observations / swath	5400	540	540	270	270	2700
	Swaths / sec	13.616	1.362	1.362	0.681	0.681	6.808
	Peak Data Rate Mbps	13.533	0.077	0.098	0.651	0.651	23.139
	% Time In Surge Mode	15%	15%	15%	15%	25%	50%
	Average Surge Data Rate Mbps	2.030	0.012	0.015	0.098	0.163	11.570
SURGE+NOMI	Data Rate per Instrument Mbps	4.906	0.014	0.018	0.186	0.241	11.575
	Data Rate per Satellite Mbps *	3.20				0.15	7.23
	TDRS link required	Single Access				Multiple Access	Single Access
	Total Data Rate By Satellite Class Mbps *	42				1.2	94
	Total Data From All LEO Satellites Mbps *	137					

* Note: Includes 2X Compression plus 25% overhead for forward error correction code and formatting

3.7.4. High Data Rate Spacecraft/Sensors (KSA)

The data rates from the Imager/Sounder satellites and the RADAR satellites are high enough that they must use the KSA service. At any one time, there could be 8 or 9 satellites in view of each TDRS and they each would have to cycle through all these 8 or 9 satellites every 15 minutes. Given that 0.5 minutes is lost between LEO satellite contacts, one SA antenna would lose 4.5 minutes every 15 minutes while slewing among 9 satellites. This leaves only 10.5 minutes for data collection from 9 satellites giving each LEO satellite only 1.1 minutes to transmit its previous 15 minutes of data. To accommodate some higher than average surges of data, we use a contact time of 1 minute or a data speed up of 15X. This would increase the data rate from the Imager/Sounder satellites to 48 Mbps and from the RADAR satellites to 109 Mbps. These data rates are well within today's communication capability.

3.7.5. Sensitivity to the Selected Data Parameters

It is believed that the parameters of Table 3-1 can be achieved with appropriate technology investment.

The number of each type of satellite and its swath width is driven by the time to achieve global coverage requirement. If the 2700 km swaths assumed in Table 3-1 cannot be achieved, more satellites will be required. The total data to be communicated will not change significantly, but the additional satellites will cause more wasted dead time for the KSA Service.

The data rates shown in Table 3-1 assume that the instruments continuously sample. This is appropriate when the satellites are in tropical latitudes, but the earth's polar regions would be oversampled. It is expected that the sampling in the polar regions would be reduced, but this was not factored into the data rate calculations because there can be 15 minute periods when most of the satellites will be in tropical latitudes.

Lossless compression of a factor of 2 was used in Table 3-1. The compression that will be achievable may be more or less than this number. However, it is seen that even if no compression were possible, the satellite data rates would still be within a TDRS's capability.

Table 3-1 assumed that the imager and sounder instruments were flown on the same satellite. If they were flown on separate satellites, the EO Imager would still require KSA service, but the others could be handled on MA. If flown on separate satellites, they would then likely fly in a close formation to achieve nearly simultaneous area coverage. In this case, while the EO Imager data was being transmitted on KSA, the data from the other three satellites could be sequentially transmitted using the SSA capability.

4. Modeling and Data Assimilation Gaps

4.1. Computing Technology

Even today, numerical weather prediction is one of the most computationally taxing functions performed. Indeed, many of the current limitations in weather prediction are imposed not by uncertainties in the science, but rather by the inability to perform the necessary calculations in time. Although advances in computer technology will lead to faster computers, the needs of the future weather forecasting architecture will also increase tremendously.

4.1.1. Anticipated Computing Technology Needs

Clearly, the future computing needs of a weather forecasting system will increase as the architecture discussed in the main part of this study report comes to fruition. A quantitative assessment identifies several aspects that will greatly impact the ultimate computational complexity of this future system. Three key elements driving the increase in computational needs are related to:

- Increases in the resolution of the analysis and model functions
- Increased complexity of algorithms contained within these functions
- Increased numbers of observational data collections providing an input to the models

4.1.1.1. Resolution Increases

As the model resolution increases by some factor in the horizontal, the number of calculations required increases by the square of that factor. A quick, qualitative assessment concludes that one would expect huge increases in processing needs as the analysis and model resolutions go from the current 1x1° globally (about 111 km resolution) to a resolution of 25 km or better. When the increases in the number of vertical layers represented by the models is also considered, these increases become even greater.

4.1.1.2. Algorithm Complexity Increases

Analysis Complexity

Current analysis schemes used by various agencies range from 3-D variational (3Dvar) analyses, spectral/statistical interpolation, or other variations of optimum interpolation schemes. The near-term future of analyses will likely progress to 4-D variational analyses, where observations are brought into the model by analyzing the data with respect to time in addition to the 3 spatial dimensions. Of course, the additional analysis dimension adds considerably to computing complexity. The computational complexity of these schemes tends to scale as the square of the number of observational data points being brought into the analysis. Although the future will likely cause an increase in observational data of two to three orders of magnitude, these increases can be offset conducting analyses at more frequent intervals (decreasing the numbers of observations ingested at each step).

Further into the future, analysis schemes will likely include Kalman filtering as a principle component. Although well-understood, Kalman filtering techniques for global analyses are computationally very expensive – prohibitively so at present. The complexity of these schemes scales as the number of observational points times the square of the number of model gridpoints. Again, the number of observations processed for a given analysis can be decreased by more frequent analyses, however, the number of model gridpoints cannot be easily reduced.

Model Complexity

The complexity of atmospheric models is subject to great variability. Because the resolution of today's models is such that certain smaller-scale features cannot be accurately modeled, they are parameterized. In many cases, the complexity of these parameterizations is greater than the explicit modeling. As the resolution of models increases, however, these features could be modeled explicitly, perhaps bringing a general decrease in the actual complexity of the NWP algorithms. However, with increasing resolution comes a greater number of model gridpoints. To what extent the competing effects will weigh is uncertain, although it is almost certain that the general trend will be towards overall increases in complexity. A linear increase with the number of modeling points does not seem unreasonable.

4.1.1.3. Observational Data Increases

With today's satellite remote sensing, many more observational data points are available than there were even 10 years ago. Typical estimates for the number of observations used by today's models center around approximate 10^6 observations per day. Even by the most conservative estimates, this number will increase by two orders of magnitude by 2025. Even with larger numbers, however, it seems reasonable to expect that no more than 10^8 observations per day will be used by the models once redundant and/or low-quality data are filtered out. Still, with analysis complexity scaling as the square of the number of observations, this will result in a huge increase in computational costs.

So far, much of the discussion has centered on qualitative assessments of the increases in computational complexity. In order to obtain a quantitative assessment, it is necessary to use estimates of future model specifications (resolution, numbers of observations, etc.) to calculate model complexity. These numbers can then be used to determine how much computing capability will be needed.

Many of the calculations used to determine the computational costs of future systems were based on Lyster, July 2000. In this paper, Lyster presents a methodology for calculating the complexity of various analysis and modeling schemes based on specifications such as number of analysis points, number of observations, time step, etc. The results of the Lyster calculations are a total number of floating point operations needed to perform the stated function. Based on assumptions of the amount of time needed to complete a given task and an estimated computational efficiency, an estimate of the required sustained computing power is obtained.

Complexity Calculator

The calculations described by Lyster have been included in a simple spreadsheet. By changing forecast system specifications (such as analysis/model resolution, number of observations, runtime, etc., an estimate of computing resources in GigaFLOPS (10^9 Floating Point Operations per second) is returned for various computational algorithms.

Input Variables

The key input variables for the calculations (and their initial values) are shown below:

- Horizontal resolution (25 km)
- Number of analysis levels (100)
- Number of model levels (100)
- Analysis interval (1 hour)
- Number of upper air analysis/prognostic variables (4)
- Number of surface analysis/prognostic variables (1)
- Number of observations per day (10^8)
- Analysis run time (10 minutes)
- Quality Control run time (10 minutes)
- 24 hour forecast run time (5 minutes)
- Targeted observation run time (10 minutes)

Based on these inputs, various portions of the Modeling and Data Assimilation System will require anywhere from 10^7 to 10^{13} GFLOPS of computational resources. In other words, the range of computational resources needed is 10^{16} to 10^{21} Floating Point Operations per Second. For the curious, the range can also be stated as 10 PetaFLOPS to 1 ZettaFLOPS.

4.1.2. Anticipated Computing Technology Capabilities

At first glance, the numbers discussed in the previous section appear so high as to be impossibly ludicrous. However, with the expected growth in computing capabilities, the lower end of this spectrum actually falls within the domain of possibility.

When Gordon Moore first observed the growth in transistor density on computer chips, he found that it doubled roughly every 18 months. Although he was not necessarily referring to computing speed, transistor density typically relates linearly to it. Thus, the assessment of a doubling of computing speed every 18 months is now widely known as Moore's Law.

There is some concern that today's conventional computing systems (e.g., silicon or CMOS-type chips) will reach a size barrier in anywhere from 15 to 20 years. However, if one extrapolates computing speed back in time before solid state computing, it becomes apparent that the computing speeds of the earlier tube computers is consistent with Moore's Law. Thus, it is not unreasonable (at least for this study) to assume that some future technology (e.g., optical or quantum computing) is likely to pick up where silicon leaves off. Thus, the future computing capabilities expected in 2025 (over the course of normal evolution) are based on application of Moore's Law to some current computing capability.

Today's state of the practice systems boast speeds on the order of hundreds of GFLOPS. Although systems capable of sustained speeds in the TFLOPS range are in use, they are still considered state of the art and not readily available for operational centers. For this study, a current capability of 500 GFLOPS was used as a baseline for application of Moore's Law. Projecting forward, a sustained computing speed in the range of $10^7 – 10^8$ GFLOPS was obtained.

4.1.3. Gap Analysis

With expected computing capabilities and a tool to estimate computing resource needs, we can now examine the technology gaps.

4.1.3.1. Technology Shortfalls

Figure 3-1 illustrates the expected growth of computing capabilities (red line) over time against the estimated computational needs of the future weather forecast system with specifications as stated in Section 4.1.1.

As discussed in the main body of the study report , the MDAS consists of several functions. These include quality control of the input data, analysis of the observed data onto a regular grid, the global forecast model, and the targeted observation selection. The spreadsheet tool estimates computation complexity of each of these functions. For the most part, the Kalman filtering analysis function is the most computationally expensive part of the future system. To explore options to reduce the system's computational requirements, several additional analysis schemes were also explored. The four threshold lines (in blue) indicate the computational needs (in GFLOPS) for the full MDAS obtained for each of these analysis options.

The topmost line indicates the processing needs for an analysis and forecast system using a full Kalman filtering analysis scheme. Because the computational needs of this algorithm scale as the square of the analysis gridpoints times the number of observations, this turns out to be computationally very expensive. This analysis places the computing needs at more than five orders of magnitude greater than the expected capabilities.

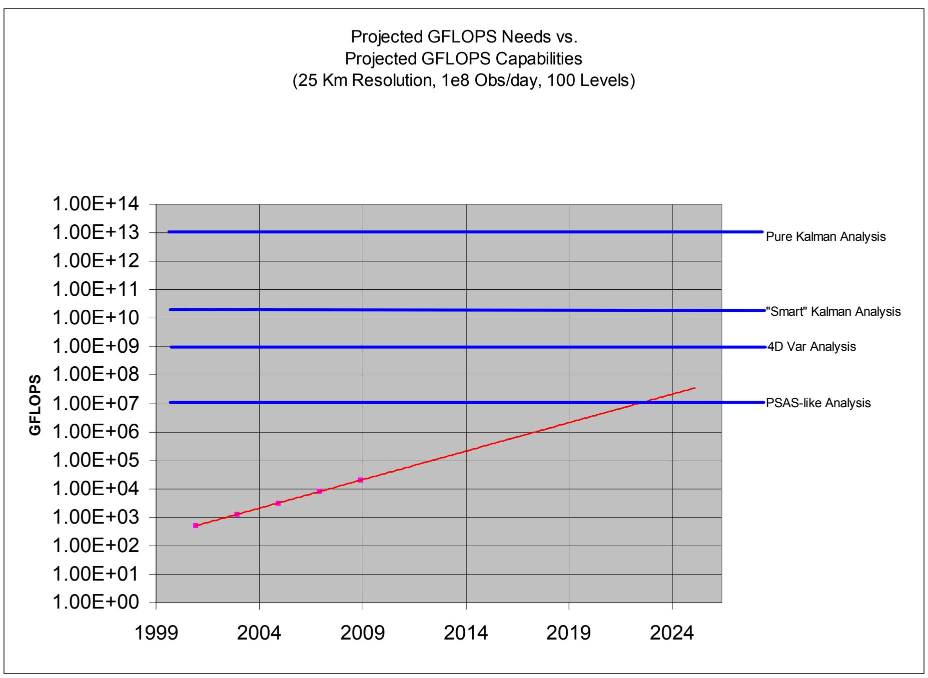

Figure 4-1 Projected Computing Capability Gap

4.1.3.2. Trade Areas

Obviously, a full Kalman filtering is most likely well beyond the capabilities of projected future technologies. However, discussions with analysis experts from DAO have led to other options. While a full Kalman filter would provide the best analysis, a partial or "smart" Kalman filter might yield a solution that is "good enough." Such an analysis might scale as the number of gridpoints raised to the power of 1.6 or 1.7 (vice 2.0). Such a scheme would realize great savings in computational costs. The value of this factor would be related to how well the Kalman analysis performs. Thus, the next lower line (labeled "Smart" Kalman Analysis) marks the threshold for a factor of 1.7. Although nearly 2½ orders of magnitude higher than expected capabilities, it is still significantly lower than a full Kalman analysis. Using a factor of 1.6, the requirements drop to 4×10^9 GFLOPS, less than 1½ orders of magnitude above the projected capabilities.

Another scheme that has a lower computational costs is the 4 dimensional variational (4D Var) analysis. As of now, 4D Var analysis schemes are seeing use in either small scale analysis of all observational data types or for global analysis of limited data types. As the implementation of this scheme improves and resources allow for the somewhat high computing costs, 4D Var will most likely see use as a full global analysis scheme in the near future (at current spatial resolutions). The estimates of the 4D Var computational costs were made for the same specifications as the Kalman filtering scheme and are indicated by the third blue line. The costs, while lower than all the previous analysis schemes, are still 1½ orders of magnitude above the expected capabilities. Additionally, the computing cost of the analysis portion of the MDAS is now within the realm of the other functions within the system, especially the global forecast model.

The final analysis scheme examined to lower computational costs is a 3 dimensional Variational (3D Var) analysis. A simpler version of the 4D Var, the complexity of this scheme is on par with today's analysis tools used operationally. When the complexity of the 3D Var is estimated using the same specifications (lower blue line on Figure 4-1), computing needs are found to be on the order of 10^7 GFLOPS – well within the projected capabilities of future computing systems. Furthermore, the computing costs of the analysis are found to be of the same order of magnitude as those of the global forecast model.

4.1.3.3. Future Technology Needs

From this examination, we see a large gap between expected computing needs and resources. There are two ways to close the gap – raise the available computing speed or lower the computational requirements. Raising computing speed can only be done through the technology advances of the computer industry. Reducing the computational requirements, however, is within the realm of the earth science community by way of computation that is more efficient or smarter data analysis algorithms.

4.1.3.4. Recommendations

ESTO should keep an eye on the computer industry and examine technologies that could lead to the capabilities needed, such as reconfigurable computers. Furthermore, NASA's earth science community should maintain an open dialog with the computer research community so they remain aware of the future computational needs

ESTO should also support research into developing more efficient computational systems and algorithms to make better use of the available computational resources. As was noted earlier, current weather codes have an efficiency of perhaps 10%. Significantly increasing this efficiency would be a good start at closing the gap. In addition to efficiency, ESTO should support research into smarter analysis and forecast algorithms such as the "smart" Kalman filter. Although related to the science of numerical weather prediction, these are actually computing technology advances that need to come about in order to close the anticipated gaps.

4.2. Meteorological Science

In addition to the technologies discussed above, significant development will have to occur in key areas of the numerical modeling arena in order to support the concepts discussed in the study report. These areas include selection of targeted observations, continuous assimilation and model self-assessment.

5. Areas for Further Study

As this study was conducted, certain assumptions were made in order to complete the study in a limited timeframe. The following sections discuss several areas in which additional study would be beneficial.

5.1. Study Detail Refinements

5.1.1. On Board Processing Trades

In the current version of the notional architecture described in the study report, only minimal data processing is accomplished on board the spacecraft conducting remote sensing measurements. Instead, the calibrated, earth located data are transmitted to the ground to be reduced for ingest into the analysis models.

It is conceivable that some portion of the data processing should be done on the spacecraft that would benefit either the efficiency of the system or quality of the forecast product. For instance, data that are reduced from the raw (sensor-based) measurements to the required parameters and spatial resolution needed for the model might be far less voluminous, decreasing the bandwidth requirements needed for downlink. As another example, data that are reduced to their desired environmental parameters might be better suited for the automated event detection needed for a rapid reconfiguring of the SensorWeb.

However, moving processing from the ground to the space platform entails its own difficulties. Some reduction schemes require supplemental data that will have to be uplinked to the spacecraft -- will the potential bandwidth savings and quality improvements be worth the additional uplink? Additionally, data reduction schemes could be computationally expensive -- can these tasks be accomplished on board with expected spacecraft capabilities? Would the benefits be worth the costs of providing these capabilities?

This portion of the follow-on study would identify several key areas where processing could be moved to the spacecraft, develop a concept of operations for the processing, and discuss the benefits versus costs of these changes.

As an example, temperature and humidity data are currently derived from infrared and microwave radiances transmitted to the ground. If the data reduction were to be performed in space, the amounts of data downlinked could be greatly reduced. However, some analysis schemes are being used that directly ingest these measured radiances rather than the derived environmental parameters. Would eliminating the availability of radiance data at the ground adversely impact the quality of the analysis? Even now, some organizations are moving away from radiance assimilation.

Other areas of data processing to be considered could include (but are not necessarily be limited to):

- Data Quality Control
- Rapid Event Detection
- Calculation of the forward model results (needed for QC)
- Data Analysis (reduction of sensor-based coordinates to model grid coordinates)

5.1.2. Assimilation and Forecast CONOPS

In the original gap study, computational requirements were estimated using an assumed concept of operations (CONOPS) for data assimilation and global forecast generation. Among the variables for which values were chosen (and that could have large impacts on the MDAS computational resource requirements) are:

- Frequency of assimilation runs: Hourly was chosen, but other intervals might produce better quality products -- even a continuous assimilation process has been suggested in the community.
- Number and types of ensemble forecasts: For current long-range forecasts, numbers of ensemble members range from 10s to 100s. Furthermore, other ensemble approaches (e.g., Monte Carlo suites) have been suggested that could greatly impact the computing resources required.
- Targeted observation methodologies: Currently proposed techniques for selecting observations to be collected revolve around calculating the adjoint of the models. Other techniques might be available that would have different levels of computational cost and produce better results.

The follow-on study should gather information from domain experts and literature review to generate options for an assimilation and forecast CONOPS. Using this updated information, computational resource requirements will be re-estimated to provide a more accurate range of values needed.

5.1.3. SensorWeb Management and Monitoring

Much of the intelligence surrounding the management and monitoring of the SensorWeb was not fully developed in the original study. Some of the aspects not fully explored include:

- Architecture requirements (e.g., communications and computing needs)
- Timing requirements for SensorWeb responses (e.g., how quickly does the SensorWeb need to respond to events detected by other portions of the SensorWeb? How does communications latency affect the response?)

The follow-on study should develop a CONOPS specific to the SensorWeb that will address the intelligence required (both distributed and integrated), define various options for the location of various portions of the intelligence, and examine how these options will affect the architecture needs. Updates to these needs will be examined to see how they might affect the gap analysis.

5.1.4. Architecture Management and Monitoring

In the original gap study, some aspects of the overall system monitoring and management were given cursory examination. This portion of the architecture is responsible for monitoring the performance of the system as a whole. These functions also provide for such things as the setting of "policy" items (e.g., forecast production schedules), approval of science community requests for data collection, or human update/intervention into system operations.

Although these functions were discussed in the study report, very little detail was provided about what impacts to the overall architecture these functions might have. The follow-on study should provide additional detail of the overall architecture.

5.2. Three to Five Day Forecast Study

This follow-on study would build on the same high level system architecture concept that emerged from this study, that was non-specific in some areas since it necessarily involved [educated] speculation on almost every relevant future technological capability from constellation management, to computing technologies, to communications. In order to bore down into the deeper meaning of the two way interaction, we think it is essential now to hold some variables constant so that we may focus on the system architecture question in more concrete terms. This increase in detail can be obtained by focusing on a well-controlled scenario that is known to be tractable (1-5 day forecasting).

For control, this new study would start by assuming only the capabilities of research and operational space-based observing systems that are being planned now for deployment in the 2010 – 2020 time frame (e.g. GPM, NPOESS, GIFTS), and/or technologies that are fully expected to have reached a prototype demonstration stage of maturity the 2010 – 2025 time frame. The basic observing characteristics of these future systems are more or less given. Also, to sustain focus on the system architecture, this follow-on study should use as science scenarios a finite set of weather phenomena whose evolution and prediction would be encompassed over time scales ranging from 24 hours to 120 hours. For example, localized severe thunderstorm forecast 24 hours in advance, prediction of East Coast snowstorms 4-5 days in advance, or prediction of devastating winter low-pressure systems that impact California in El Nino years. By naming the scenarios and phenomena of interest up front, and knowing the class of observations that will or should be available, a higher level of specificity can be obtained. This will permit a better focus on how the entire system must operate and be designed in order to provide the needed coordination between and among space and terrestrial based observing systems, and operational weather modeling systems.

Given more concrete notions based on realistic use-case scenarios of the data flows and desired interactions among the system components, it will be possible then to consider the system logic, architectures and technologies, as well as advances in system theory, communications and that could provide the necessary interactivity and results from a highly intelligent, highly integrated operational weather forecast system not only for short to medium range forecasting but out to 10 – 14 days.